

LA-UR-21-21979

Approved for public release; distribution is unlimited.

Title: Porting E3SM from LANL's Open to Classified Computing Networks

Author(s): Price, Stephen F. Dr
Harp, Dylan Robert
D'Angelo, Gennaro
Green, Jennifer Kathleen
Wilson, Cathy Jean

Intended for: Report

Issued: 2021-02-26

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Porting E3SM from LANL's Open to Classified Computing Networks

S.F. Price (T-3), D. Harp (EES-16), G. D'Angelo (T-3), J. Green (HPC-ENV), C. Wilson (EES-14)

Summary

This project used turquoise computing time to support a LANL TED project focused on porting E3SM to classified computing resources. Computing time was used to test and confirm a complete and working simulation and analysis workflow prior to porting to yellow and SCI-level computing platforms. The code port and all planned simulations and analysis were successful. E3SM analysis from simulations on classified platforms was published to C2S and will be used to demonstrate and advertise LANL's classified Earth system modeling and analysis capabilities to current and new customers in the intelligence community. This project also significantly improved library and E3SM analysis support on turquoise, yellow, and classified computing platforms at LANL, which will result in benefits beyond the work discussed here.

Overview

This project requested a 2020 LANL HPC allocation in support of a LANL Technology and Demonstration (TED) project that focused on porting DOE's Energy Exascale Earth System Model (E3SM) to classified LANL HPC resources (Sage). 23K node hours on *Grizzly* (awarded under allocation "w20_e3sm2sage") was used to run (1) decadal-scale simulations with E3SM in standard, low-resolution and new, variable resolution configurations, (2) add software support for and conduct new analysis of these same simulations, and (3) insure a complete "end-to-end" simulation and analysis workflow in an open computing environment (where links to external code and data are allowed / automated and where a wide range of project expertise is available) to a highly classified, closed computing environment (where external links are explicitly cut and staffing expertise is extremely limited).

Scientifically, this effort focused on the initial testing and analysis of a new, variable-resolution E3SM model configuration – an Arctic Regionally Refined Mesh (ARRM) – with respect to its fidelity in simulating a realistic Arctic climate. The resulting simulations and analysis conducted on *Grizzly* helped identify a number of model parameterization and grid resolution choices that require further improvement and testing by the broader E3SM teams at LANL, and that will be incorporated into ongoing simulation campaigns under the E3SM and HiLAT projects.

After a complete and successful workflow was first demonstrated on *Grizzly*, all source code, third-party libraries, and input and analysis datasets were ported to *Fog*, where a limited set of simulations and analysis were repeated. This demonstrated an initial successful port to a computing environment where external code links could be severed for testing purposes. Finally, the port was repeated but to LANL's *Sage* cluster, where a number of production simulations and analysis were completed. Select analysis of these production simulations was chosen for publication to C2S, where it will be used to demonstrate and advertise LANL's classified Earth system modeling and analysis capabilities to current and new customers in the intelligence community. These TED project deliverables thus contribute to LANL's broader efforts at advancing its Climate Impacts on National Security (CINS) capabilities. In particular, with the change in the US administration, we have received interest from the National Intelligence Council

about this unique capability, and we will be including materials from this project in briefings to a broad range of agencies in March 2021.

In addition to the planned project outputs noted above, several unplanned but substantial benefits came from this work. First, through an unplanned collaboration with LANL HPC staff, a new and updated software stack was setup using *Spack*, ultimately supporting 5 E3SM model build configurations, using two compilers (gnu and intel) and three MPI libraries (openmpi, mvapich2, and impi). This software stack, first assembled and tested on *Grizzly*, was ported to both *Fog* and *Sage* and represents a substantial improvement in flexibility and robustness for supporting E3SM on LANL HPC platforms (Previously, LANL staff on the E3SM project maintained their own library support in an *ad hoc*, fragile, and time-consuming process). Second, HPC staff also helped to get *Paraview* – already well supported on LANL’s open computing platforms – available and working on the HAL and HAL0 SCI classified computing networks. This industry standard HPC analysis tool is now available to any computing projects utilizing SCI-level clusters and networks at LANL.

Detailed list of project outputs

- Improved software library support for the E3SM project on *Grizzly*, *Fog*, and *Sage*
- Updated E3SM source code and support for new grid configurations on *Grizzly*, *Fog*, and *Sage*
- Updated and expanded E3SM analysis capabilities on *Grizzly*, *Fog*, and *Sage* (specifically, *MPAS-Analysis* software, used for ocean and ice model analysis, was updated and *E3SM-diags* software, used for atmosphere and land model analysis, was added)
- Added support for *Paraview* HPC analysis software to LANL’s SCI-level computing networks (HAL and HAL0)
- Completed simulation and analysis of 375 E3SM model years on *Sage* (16.5 M cpu hrs, 53 Tb of model output)
- Use of ~85% of *Sage* for high resolution E3SM configuration with peak model throughput of ~5 simulated model years per day
- Published select simulation analysis products to C2S for demonstrating LANL classified Earth system modeling and analysis capabilities to current and new customers in the national security and intelligence communities

Additional information on project outputs is included in Part 2 of this report, in the form of several Powerpoint slides.